



Enhanced Media Access Control with Echelon's LONTALK™ Protocol

August 1991

LONWORKS™ Engineering Bulletin

Introduction

This note provides an introduction to the LONTALK™ media access control (MAC) sublayer. The MAC sublayer is part of the Data Layer of the OSI Reference Model. Many different MAC algorithms exist in networks today. One family of these algorithms is called CSMA (Carrier Sense Multiple Access). The MAC algorithm used by the LONTALK protocol belongs to the CSMA family. An explanation of the CSMA algorithm used by the LONTALK protocol is presented below. Further, its features are compared with some of the other members of the CSMA family.

Background on CSMA Family

The CSMA family of media access control algorithms requires a node to establish that the medium is idle before it begins to transmit. However, each algorithm behaves differently once the idle state is detected. This results in very different network performance results under conditions of heavy data traffic.

Some CSMA algorithms use discrete intervals of time called slots, or randomizing slots, to access the medium. By limiting access to the medium by a given node to specific time slots, slotted media access greatly reduces the probability of two packets colliding. Slotted media access is used by p-persistent CSMA, and Echelon's LONTALK CSMA protocol.

Overview of the LONTALK MAC Sublayer

Echelon's LONTALK protocol uses a new CSMA MAC algorithm called *predictive p-persistent* CSMA. The LONTALK protocol retains the benefits of CSMA but overcomes its shortcomings for control applications. Existing media access control algorithms such as IEEE 802.2, 802.3, 802.4, and 802.5 do not meet all the LONTALK requirements for multiple communication media, sustained performance during heavy loads, and support for large networks.

As in p-persistent CSMA, all LONWORKS nodes randomize their access to the medium. This avoids the otherwise inevitable collision that results when two or more nodes are waiting for the network to go idle so that they can send a packet. If they wait for the same duration after backoff and before retry, repeated collisions will result. Randomizing the access delay reduces collisions. In the LONTALK protocol, nodes randomize over a minimum of 16 different levels of delay called randomizing slots. Thus the average delay in an idle network is eight slot widths.

In p-persistent CSMA when a node has a message to send, it does so in a given randomizing slot with probability p . However, the LONTALK protocol carries the added improvement that p is dynamically adjusted based upon network load. When the network is idle, all nodes randomize over only 16 slots. When the estimated network load increases, nodes may randomize a large number of slots. The number of slots increases by a factor of n , where the range of n is from 1 to 63. Echelon calls n the estimated channel backlog. It represents the number of nodes with a packet to send during the next packet cycle.

This method of estimating the backlog and dynamically adjusting the media access allows the LONTALK protocol to have just a few randomizing slots during periods of light load, while having the benefit of many randomizing slots during periods of heavy load. Thus, media access delays are minimized during periods of light load, and collisions are minimized during periods of heavy load.

As noted earlier, the ability to adjust the number of randomizing slots depends on the ability to estimate the channel backlog. In the LONTALK protocol, a transmitting node includes information in the packet on the number of acknowledgements expected as a result of sending that packet. All the nodes that receive the packet increment the estimated channel backlog by that amount. Likewise, the estimated channel backlog is decremented by 1 at the end of each packet cycle. The estimated channel backlog is never decremented below 1. Since LONTALK packets are typically acknowledged, 50% or more of the channel backlog is predictable at any time.

Benefits of LONTALK MAC Sublayer

LONWORKS systems allow thousands of nodes and multiple media on a single network. Because of the characteristics of different communication media and the potential need to cover large distances, LONWORKS networks must be able to support low data rates. Time-slotted MAC sublayer s are not appropriate for low data rates because they are suited to high data rates and small numbers of nodes. Thus, a time-slotted MAC sublayer was not chosen for the LONTALK protocol.

The multiple-media support provided by the LONTALK protocol also rules out a token-ring approach which only works on media with the propagation characteristics of wire. In a token-ring network, the token passes in an orderly fashion around the ring. This cannot work on either powerline or RF networks because all stations receive the token simultaneously. Additionally, the cost of incremental hardware to recover the token when it is lost, and to rapidly acknowledge the token, makes the hardware for this approach more expensive to implement.

A token bus architecture solves the problem of sequential passing of the token by including addressing information in the token. However, at low data rates, the process of circulating the token can result in considerable token latency. Since a node cannot transmit without first possessing the token, this latency adversely

affects response time. Additionally, token bus systems must reconfigure themselves each time a new device either becomes active or drops out. This overhead to reconfigure is a problem for all token bus networks. Since reconfiguration brings the network down for its duration, battery-powered nodes whose normal operation is to wake up, send some messages, and power down, would cause the token bus system to suffer frequent reconfigurations. Battery-powered nodes are required for applications needing RF or IR communication, for security applications, and for fire/life safety applications, to name a few.

The CSMA family of MAC sublayers does not require a ring topology, synchronization or reconfiguration, and does permit nodes to drop out and rejoin the network transparently. Additionally, it supports many nodes and is inexpensive to implement in hardware. Unfortunately, CSMA/CD (for example IEEE 802.3) behaves poorly during periods of overload, so it is generally not used for control applications. P-persistent CSMA works very well for small values of p at the expense of additional delay during relatively idle periods. The LONTALK MAC sublayer has the advantages of p-persistent CSMA without the disadvantages of additional delay during low traffic, or significantly reduced throughput under conditions of high traffic.

In summary, the LONTALK MAC sublayer specifically overcomes the shortcomings of existing MAC sublayers in the following areas: multiple-media communication, low data rates, sustained performance during conditions of heavy network traffic, and large networks.

The LONTALK Predictive P-Persistent CSMA Protocol

As mentioned earlier, Echelon's LONTALK protocol falls under the category of predictive p-persistent CSMA. It is predictive because each node dynamically predicts how many other nodes have a packet to send at any given time. This prediction influences the number of randomizing slots between each packet. The higher the prediction, the more slots there are for the nodes to randomize over. Increasing the number of slots reduces the probability of a collision. The predictive algorithm is based upon the fact that most LONTALK packets are acknowledged. The number of acknowledgements that a given packet generates is encoded into the packet. Each node on the channel receives the packet and adds the number of acknowledgements to the channel backlog. If none of the packets is acknowledged, the predictive part of the algorithm does not dynamically expand the number of randomizing slots with an increase in load. Using exclusively unacknowledged services causes the LONTALK protocol to behave like a p-persistent CSMA where $p = 0.0625$. However, p-persistent CSMA is still significantly better than IEEE 802.3 under conditions of heavy network traffic.




Figure 1: Effective network throughput versus offered traffic on a LONWORKS network using unacknowledged services. Note that in this mode, LONTALK behaves like conventional, p-persistent CSMA.

The results of two experiments used to illustrate this point are shown in figures 1 and 2. The graphs shown in these figures illustrate results achieved from a test bed of 24 nodes. Twenty-two of the nodes acted as traffic generators. The other two acted as test nodes, repeatedly sending messages to each other at different traffic levels. Both graphs show the amount of actual traffic that got through on the network versus the offered traffic, or the number of packets attempted for transmission.

Figure 1 shows this data using exclusively *unacknowledged* packets. In this mode, the network behavior is similar to a 0.0625-persistent CSMA algorithm. Note how the network throughput rises rapidly with offered traffic and plateaus at 300 packets/sec. The number of failed transactions jumps at the point of network saturation (network throughput at 375 packets/sec) and continues to rise. At the same time, network throughput starts to fall off due to excessive collisions when the network is driven well beyond saturation.

To demonstrate the effectiveness of the predictive algorithm, the experiment was re-run with only one change. In this case, messages used *acknowledged* services. Figure 2 shows the data derived from this experiment. Note in this graph that there were no instances of messages failing to get through (with a retry count of 3). Note also that network throughput does not degrade with an increase in offered traffic

past the saturation point. Additionally, it is important to note that the collision rate levels off rather than continuing to increase with offered traffic. This active management of the collision rate is what makes the LONTALK protocol's predictive p-persistent CSMA algorithm superior to other CSMA algorithms.




Figure 2: Effective network throughput versus offered traffic on a LONWORKS network using acknowledged services. Note that in this mode, predictive p-persistent CSMA allows active management of the collision rate and prevents throughput from degrading past the saturation point.

Conclusion

The LONTALK protocol uses a new CSMA MAC algorithm called predictive CSMA. To avoid collisions, all LONWORKS nodes randomize their access to the communication medium using time slots. The LONTALK protocol dynamically adjusts the number of randomizing time slots by predicting the channel backlog. By actively managing the collision rate, the LONTALK protocol provides a superior MAC sublayer for multiple-media communication, low data rates, sustained performance during heavy loads, and large networks.

Disclaimer

Echelon Corporation assumes no responsibility for any errors contained herein.
No part of this document may be reproduced, translated, or transmitted in any form without permission from Echelon.

© 1991 Echelon Corporation. ECHELON, LON, and NEURON are U.S. registered trademarks of Echelon Corporation. LONMANAGER, LONBUILDER, LONTALK, LONWORKS, 3150, and 3120 are trademarks of Echelon Corporation. Patented products. Other names may be trademarks of their respective companies. Some of the LONWORKS TOOLS are subject to certain Terms and Conditions. For a complete explanation of these Terms and Conditions, please call 1-800-258-4LON.

Echelon Corporation
4015 Miranda Avenue
Palo Alto, CA 94304
Telephone (415) 855-7400
Fax (415) 856-6153

Echelon Europe Ltd
105 Heath Street
London NW3 6SS
England
Telephone (071) 431-1600
Fax (071) 794-0532
International Telephone + 44 71 431-1600
International Fax + 44 71 794-0532

Echelon Japan K.K.
AIOS Gotanda Building #808
10-7, Higashi-Gotanda 1-chome,
Shinagawa-ku, Tokyo 141, Japan
Telephone (03) 3440-8638
Fax (03) 3440-8639